

## D E S C R I P T I O N

## METHOD AND SYSTEM FOR PROCESSING DATA PACKETS

## BACKGROUND OF THE INVENTION

## 1. Field of the Invention

The present invention relates generally to a method and system for processing data packets in switched communications networks and in particular to a method and a system for forwarding data packets in a router.

## 2. Description of the Prior Art

A switched communications network transfers data from source to destination through a series of network nodes. Switching can be done in one of two ways. In a circuit-switched network, a dedicated connection is established through the network and is held for as long as communication is necessary. An example of this type of network is the traditional telephone system.

A packet-switched network, on the other hand, routes data in small pieces called packets, each of which proceeds independently through the network. In a process called store-and-forward, each packet is temporarily stored at each intermediate node, then forwarded when the next link becomes available. In a connection-oriented transmission scheme, each packet takes the same route through the network, and thus all packets usually arrive at the destination in the order in which they were sent. Conversely, each packet may take a different path through the network in a connectionless or datagram scheme. Since datagrams may not arrive at the destination

Parameter	Value	Unit
Temperature	25.0	°C
Pressure	1.0	atm
Flow rate	1.0	L/min
Concentration	1.0	g/L
Time	1.0	h
Volume	1.0	L
Mass	1.0	g
Energy	1.0	J
Power	1.0	W
Force	1.0	N
Distance	1.0	m
Area	1.0	m <sup>2</sup>
Volume	1.0	m <sup>3</sup>
Mass	1.0	kg
Energy	1.0	J
Power	1.0	W
Force	1.0	N
Distance	1.0	m
Area	1.0	m <sup>2</sup>
Volume	1.0	m <sup>3</sup>
Mass	1.0	kg
Energy	1.0	J
Power	1.0	W
Force	1.0	N
Distance	1.0	m
Area	1.0	m <sup>2</sup>
Volume	1.0	m <sup>3</sup>
Mass	1.0	kg
Energy	1.0	J
Power	1.0	W
Force	1.0	N
Distance	1.0	m
Area	1.0	m <sup>2</sup>
Volume	1.0	m <sup>3</sup>
Mass	1.0	kg
Energy	1.0	J
Power	1.0	W
Force	1.0	N
Distance	1.0	m
Area	1.0	m <sup>2</sup>
Volume	1.0	m <sup>3</sup>
Mass	1.0	kg
Energy	1.0	J
Power	1.0	W
Force	1.0	N
Distance	1.0	m
Area	1.0	m <sup>2</sup>
Volume	1.0	m <sup>3</sup>
Mass	1.0	kg
Energy	1.0	J
Power	1.0	W
Force	1.0	N
Distance	1.0	m
Area	1.0	m <sup>2</sup>
Volume	1.0	m <sup>3</sup>
Mass	1.0	kg
Energy	1.0	J
Power	1.0	W
Force	1.0	N
Distance	1.0	m
Area	1.0	m <sup>2</sup>
Volume	1.0	m <sup>3</sup>
Mass	1.0	kg
Energy	1.0	J
Power	1.0	W
Force	1.0	N
Distance	1.0	m
Area	1.0	m <sup>2</sup>
Volume	1.0	m <sup>3</sup>
Mass	1.0	kg
Energy	1.0	J
Power	1.0	W
Force	1.0	N
Distance	1.0	m
Area	1.0	m <sup>2</sup>
Volume	1.0	m <sup>3</sup>
Mass	1.0	kg
Energy	1.0	J
Power	1.0	W
Force	1.0	N
Distance	1.0	m
Area	1.0	m <sup>2</sup>
Volume	1.0	m <sup>3</sup>
Mass	1.0	kg
Energy	1.0	J
Power	1.0	W
Force	1.0	N
Distance	1.0	m
Area	1.0	m <sup>2</sup>
Volume	1.0	m <sup>3</sup>
Mass	1.0	kg
Energy	1.0	J
Power	1.0	W
Force	1.0	N
Distance	1.0	m
Area	1.0	m <sup>2</sup>
Volume	1.0	m <sup>3</sup>
Mass	1.0	kg
Energy	1.0	J
Power	1.0	W
Force	1.0	N
Distance	1.0	m
Area	1.0	m <sup>2</sup>
Volume	1.0	m <sup>3</sup>
Mass	1.0	kg
Energy	1.0	J
Power	1.0	W
Force	1.0	N
Distance	1.0	m
Area	1.0	m <sup>2</sup>
Volume	1.0	m <sup>3</sup>
Mass	1.0	kg
Energy	1.0	J
Power	1.0	W
Force	1.0	N
Distance	1.0	m
Area	1.0	m <sup>2</sup>
Volume	1.0	m <sup>3</sup>
Mass	1.0	kg
Energy	1.0	J
Power	1.0	W
Force	1.0	N
Distance	1.0	m
Area	1.0	m <sup>2</sup>
Volume	1.0	m <sup>3</sup>
Mass	1.0	kg
Energy	1.0	J
Power	1.0	W
Force	1.0	N
Distance	1.0	m
Area	1.0	m <sup>2</sup>

Table 1. Demographic characteristics of the study population	
Age (years)	50.0 ± 10.0
Gender	
Male	50.0%
Female	50.0%
Education (years)	12.0 ± 2.0
Income (USD/month)	1,000.0 ± 200.0
Marital status	
Married	80.0%
Single	20.0%
Divorced	0.0%
Widowed	0.0%
Occupation	
Professional	30.0%
Managerial	20.0%
Technical	10.0%
Service	20.0%
Unemployed	20.0%
Health status	
Good	70.0%
Fair	20.0%
Poor	10.0%
Chronic diseases	
Hypertension	30.0%
Diabetes	10.0%
Heart disease	20.0%
Stroke	5.0%
Other	15.0%
Smoking status	
Smoker	20.0%
Non-smoker	80.0%
Alcohol consumption	
Regular	10.0%
Occasional	20.0%
Never	70.0%

Table 1. Continued	
1.00	1.00
0.99	0.99
0.98	0.98
0.97	0.97
0.96	0.96
0.95	0.95
0.94	0.94
0.93	0.93
0.92	0.92
0.91	0.91
0.90	0.90
0.89	0.89
0.88	0.88
0.87	0.87
0.86	0.86
0.85	0.85
0.84	0.84
0.83	0.83
0.82	0.82
0.81	0.81
0.80	0.80
0.79	0.79
0.78	0.78
0.77	0.77
0.76	0.76
0.75	0.75
0.74	0.74
0.73	0.73
0.72	0.72
0.71	0.71
0.70	0.70
0.69	0.69
0.68	0.68
0.67	0.67
0.66	0.66
0.65	0.65
0.64	0.64
0.63	0.63
0.62	0.62
0.61	0.61
0.60	0.60
0.59	0.59
0.58	0.58
0.57	0.57
0.56	0.56
0.55	0.55
0.54	0.54
0.53	0.53
0.52	0.52
0.51	0.51
0.50	0.50
0.49	0.49
0.48	0.48
0.47	0.47
0.46	0.46
0.45	0.45
0.44	0.44
0.43	0.43
0.42	0.42
0.41	0.41
0.40	0.40
0.39	0.39
0.38	0.38
0.37	0.37
0.36	0.36
0.35	0.35
0.34	0.34
0.33	0.33
0.32	0.32
0.31	0.31
0.30	0.30
0.29	0.29
0.28	0.28
0.27	0.27
0.26	0.26
0.25	0.25
0.24	0.24
0.23	0.23
0.22	0.22
0.21	0.21
0.20	0.20
0.19	0.19
0.18	0.18
0.17	0.17
0.16	0.16
0.15	0.15
0.14	0.14
0.13	0.13
0.12	0.12
0.11	0.11
0.10	0.10
0.09	0.09
0.08	0.08
0.07	0.07
0.06	0.06
0.05	0.05
0.04	0.04
0.03	0.03
0.02	0.02
0.01	0.01

Table 1. Demographic characteristics of the study population	
Age (years)	Mean (SD)
18-24	20.5 (2.5)
25-34	29.5 (4.5)
35-44	39.5 (5.5)
45-54	49.5 (6.5)
55-64	59.5 (7.5)
65-74	69.5 (8.5)
75-84	79.5 (9.5)
85-94	89.5 (10.5)
95-104	99.5 (11.5)
105-114	109.5 (12.5)
115-124	119.5 (13.5)
125-134	129.5 (14.5)
135-144	139.5 (15.5)
145-154	149.5 (16.5)
155-164	159.5 (17.5)
165-174	169.5 (18.5)
175-184	179.5 (19.5)
185-194	189.5 (20.5)
195-204	199.5 (21.5)
205-214	209.5 (22.5)
215-224	219.5 (23.5)
225-234	229.5 (24.5)
235-244	239.5 (25.5)
245-254	249.5 (26.5)
255-264	259.5 (27.5)
265-274	269.5 (28.5)
275-284	279.5 (29.5)
285-294	289.5 (30.5)
295-304	299.5 (31.5)
305-314	309.5 (32.5)
315-324	319.5 (33.5)
325-334	329.5 (34.5)
335-344	339.5 (35.5)
345-354	349.5 (36.5)
355-364	359.5 (37.5)
365-374	369.5 (38.5)
375-384	379.5 (39.5)
385-394	389.5 (40.5)
395-404	399.5 (41.5)
405-414	409.5 (42.5)
415-424	419.5 (43.5)
425-434	429.5 (44.5)
435-444	439.5 (45.5)
445-454	449.5 (46.5)
455-464	459.5 (47.5)
465-474	469.5 (48.5)
475-484	479.5 (49.5)
485-494	489.5 (50.5)
495-504	499.5 (51.5)
505-514	509.5 (52.5)
515-524	519.5 (53.5)
525-534	529.5 (54.5)
535-544	539.5 (55.5)
545-554	549.5 (56.5)
555-564	559.5 (57.5)
565-574	569.5 (58.5)
575-584	579.5 (59.5)
585-594	589.5 (60.5)
595-604	599.5 (61.5)
605-614	609.5 (62.5)
615-624	619.5 (63.5)
625-634	629.5 (64.5)
635-644	639.5 (65.5)
645-654	649.5 (66.5)
655-664	659.5 (67.5)
665-674	669.5 (68.5)
675-684	679.5 (69.5)
685-694	689.5 (70.5)
695-704	699.5 (71.5)
705-714	709.5 (72.5)
715-724	719.5 (73.5)
725-734	729.5 (74.5)
735-744	739.5 (75.5)
745-754	749.5 (76.5)
755-764	759.5 (77.5)
765-774	769.5 (78.5)
775-784	779.5 (79.5)
785-794	789.5 (80.5)
795-804	799.5 (81.5)
805-814	809.5 (82.5)
815-824	819.5 (83.5)
825-834	829.5 (84.5)
835-844	839.5 (85.5)
845-854	849.5 (86.5)
855-864	859.5 (87.5)
865-874	869.5 (88.5)
875-884	879.5 (89.5)
885-894	889.5 (90.5)
895-904	899.5 (91.5)
905-914	909.5 (92.5)
915-924	919.5 (93.5)
925-934	929.5 (94.5)
935-944	939.5 (95.5)
945-954	949.5 (96.5)
955-964	959.5 (97.5)
965-974	969.5 (98.5)
975-984	979.5 (99.5)
985-994	989.5 (100.5)
995-1004	999.5 (101.5)
1005-1014	1009.5 (102.5)
1015-1024	1019.5 (103.5)
1025-1034	1029.5 (104.5)
1035-1044	1039.5 (105.5)
1045-1054	1049.5 (106.5)
1055-1064	1059.5 (107.5)
1065-1074	1069.5 (108.5)
1075-1084	1079.5 (109.5)
1085-1094	1089.5 (110.5)
1095-1104	1099.5 (111.5)
1105-1114	1109.5 (112.5)
1115-1124	1119.5 (113.5)
1125-1134	



originating input port, from where the packet gets forwarded to determined output port. In the parallel router architecture, a bottleneck or a single point of failure for the whole device might become the pool interconnect or a load balancing device of the pool.

#### SUMMARY OF THE INVENTION

It is therefore an object of the present invention to provide an improved method and system for processing data packets. The foregoing object is achieved as it is now described.

A method and a system are provided for processing data packets in a router. The router includes a plurality of input ports, a plurality of output ports and more than one packet processing units. The packet processing units derive from a piece of information associated to each data packet one of said plurality of output ports to forward the data packet to. In response to a data packet arriving at one of the input ports one packet processing unit of said multiple packet processing units is determined. The determined packet processing unit is then requested to derive a respective output port to forward the data packet to, whereby the respective output port is derived from a piece of information associated to the data packet. In the following, an identification identifying the respective output port is returned to the requesting unit. In addition, other information about the desired packet treatment and packet alterations, based on the packet processing, may also be sent back to the requesting unit. Finally, the desired treatment and alterations are applied to the data packet and the data packet is forwarded to the identified output port.

In a preferred embodiment of the method and system according to the present invention determining one packet processing unit is based on a split of an identifier vector space, where an identifier vector consists of a selected set of fields within the packet and

TOC20"0236860

The method and system according to the present invention optimize advantageously resource utilization. Furthermore, it leads to higher packet processing speed and helps to lower the costs and power requirements. Another advantage of the provided method and system is that it can cope with asymmetrical traffic load and additionally provides optimized load balancing. Furthermore, the method and system in accordance with the present invention avoid single points of failure and therefore provide fault tolerance.

### BRIEF DESCRIPTION OF THE DRAWINGS

The novel features of the invention are set forth in the appended claims. The invention itself, however, as well as a preferred mode of use, further objectives, and advantages thereof, will best be understood by reference to the following detailed description of an illustrative embodiment when read in conjunction with the accompanying drawings, wherein:

Fig. 1 is a high-level block diagram illustrating a packet processing scheme in a router being implemented in accordance with the present invention;

Fig. 2 is a high-level block diagram illustrating a preferred embodiment of a router being implemented in accordance with the present invention; and

Fig. 3 is a high-level block diagram illustrating load balancing and feedback in router being implemented in accordance with the present invention.

#### DESCRIPTION OF THE PREFERRED EMBODIMENT

With reference now to the figures and, in particular, with reference to Fig. 1, there is depicted a high-level block diagram illustrating a packet processing scheme in a router 110 being implemented in accordance with the present invention. The router 110 includes N input/output line cards LC 11 to LC N, wherein N is an integer number greater than one. The router 110 further includes M packet processing units PPU 11 to PPU M, wherein M is as well an integer number greater than one. Preferably, by far more than two line cards LC 11 to LC N and packet processing units PPU 11 to PPU M are provided within a router.

The line cards LC 11 to LC N function as an interface between a transmission facility (not shown) and a switching unit (not shown). The transmission facility forms a network through which a data packet 112 is being sent. Whereas, the switching unit transports the data packet 112 to the particular outbound line card from which it leaves the router 110. Each line card LC 11 to LC N comprises an input/output port 114 which at the same time form the input ports and output ports of the entire router 110.

The packet processing units PPU 11 to PPU M determine for a given piece of information included in each data packet 112 an outbound line card to which a particular data packet needs to get forwarded to as well as desired treatment and alterations to be applied to

the packet. For example, the packet processing units PPU 11 to PPU M convert a given destination address into a corresponding next hop physical address. This is normally performed by using a routing table containing information needed for the conversion.

The data packet 112 consists of a header portion 116 and a data portion 118. The header portion 116, or short the header 116, contains besides other information the destination address of the respective data packet 112. More detailed information about the packet is carried, for example, in the flow ID, a 5-tuple consisting of a protocol number, source and destination port and source and destination address. Whereby, the destination address is, in fact, an element of the flow ID.

The data portion 118 is formed by the data being transmitted. It is also called "payload".

Let's assume that the data packet 112 appears at the router 110 at input port 114 of line card LC 12 as indicated with arrow 120. At first, the packet 112 is parsed to extract the relevant packet fields (like, for example, the destination address). At the same time, the data packet 112 is stored until it can be forwarded.

Then, the line card LC 12 determines one of the packet processing units PPU 11 to PPU M. There are different possibilities for determining one of the packet processing units PPU 11 to PPU M according to the present invention.

In general, the computation for determining the correct processing unit takes as an input any set of fields included in the packet. Such a set of fields is referred to as an "identifier vector". For example, the identifier vector can be the flow ID (a vector consisting of five fields) or the destination address (one-dimensional vector), or any other combination of fields or

TOP SECRET

parts of fields as well. When the format of the identifier vector is defined, it forms an identifier vector space, i.e., all possible values of the vector. So, for example, in case of a vector consisting merely of the destination address, the vector space would be the address space.

In case it is preferred to preserve flows, i.e. packets having the same flow ID be mapped to the same processing unit, only such fields might be used in the identifier vector that do not change within one flow.

In case the processing units perform the longest prefix match lookup on some field, it may be preferable to define only the prefix part of the relevant field as the identifier vector. That way, prefix-defined chunks of the field range would be mapped to the same processing unit and the processing unit may be able to exploit the created traffic locality to speed up the longest prefix match lookup.

Preferably, the packet processing unit PPU 11 to PPU M is determined by computing over the identifier vector. That is, the computation is performed in such a way that to each packet processing unit PPU 11 to PPU M only packets containing identifier vectors belonging to a certain subspace of the identifier vector space are assigned to for packet processing. In other words, the identifier vector space is split into subspaces and each packet processing unit processes the requests for the identifier vector belonging to a particular subspace, whereby the identifier vector, for example, consists of the destination addresses. In a preferred embodiment the resulting split of the identifier vector space is exploited by the processor as it adjusts its processing method adaptively to the created traffic locality. Thus, a significant reduction in the number of memory accesses and, consequently, a speedup in the packet processing can be achieved. In other words,



the packet processing units PPU 11 to PPU M exploit the knowledge of the method of determining the particular packet processing unit PPU 11 to PPU M for processing the data packet in order to advantageously adjust their packet processing methods to take advantage of the said knowledge.

By using known methods from the fields of hashing and distributed caching, it is possible to provide a fully distributed scheme, i.e., all the assignment decisions can be purely deterministic and require only a few basic operations which can be computed at run-time. At the same time, these methods provide load balancing over all available packet processing units PPU 11 to PPU M. This can be achieved as described in the following.

A function  $f_1(\text{identifier vector})$  maps entries from the identifier vector space to the appropriate packet processing unit PPU 11 to PPU M. In order to take into account differences in the performance of each packet processing unit or unequal load distribution a partitioning vector  $p = (p_1, p_2, p_3, \dots, p_M)$  is advantageously introduced. The partitioning vector  $p$  defines a size of a fraction of the identifier vector space assigned to each packet processing units PPU 11 to PPU M.

Based on the partitioning vector  $p$ , a weights vector  $x$  is computed, which is then used in an extended function  $f_2(\text{identifier vector}, x)$ . The weights vector  $x$  is stored in each line card LC 11 to LC N for computing the function  $f_2(\text{identifier vector}, x)$ . The function  $f_2(\text{identifier vector}, x)$  computes the index of the packet processing unit PPU to be utilized for a specific data packet as a function of the identifier vector and the weights vector  $x$ .

The splitting of the identifier vector space is advantageously performed by a method being an adaptation of a method called Highest Random Weight method (HRW) that is described in D. G.

0969370-07001  
T03040-04235960

Thaler, C.V. Ravishankar - "Using Name Based Mappings to Increase Hit Rates", IEEE/ACM Transactions on Networking, Vol. 6, No. 1, February 1998 or K. Ross - "Hash-Routing for Collections of Shared Web Caches", IEEE Network, Vol. 11, No. 6, November - December 1997.

The split of the identifier vector space is determined by assigning to each packet processing unit PPU 11 to PPU M a numerical quantity. Whereby, the numerical quantity results of a multiplication of a pseudorandom function  $\text{rand}()$  and a weights factor  $x_j$ ,  $x_i$  taken from the weights vector  $x$ . Parameters of the pseudorandom function  $\text{rand}()$  are identifiers  $i$ ,  $j$  indicating a particular packet processing unit, and the identifier vector of the packet to be processed. Furthermore, the result of the pseudorandom function  $\text{rand}()$  is multiplied with the according element  $x_j$  or  $x_i$  of the weights vector  $x$ , respectively. Then, the packet processing unit is selected which has the highest numerical quantity assigned to it.

$$f2(\text{identifier vector}, x) = j \Leftrightarrow x_j \cdot \text{rand}(\text{identifier vector}, j) = \max x_i \cdot \text{rand}(\text{identifier vector}, i) \text{ over all packet processing units } i$$

This scheme is fully distributed, has low overhead and provides load balancing and minimal disruption in case of remapping, when one or more packet processing units PPU 11 to PPU M fail and the workload has to be reassigned to the remaining ones. This function also takes into account different processing capacities of the packet processing units represented by the weights vector  $x$ .

This scheme can also be coupled with the Fibonacci hashing scrambling method, which leads to uniformly distributed sequences, such a mapping scheme can very simply be implemented. The Fibonacci hashing method is descibed, e.g., in D. E. Knuth „The Art of

T0E0/0"0288660

Computer Programming, Vol. 3, Sorting and Searching", Addison - Wesley, 1973.

After determining one of the packet processing units PPU 11 to PPU M in one of the aforementioned ways a request is sent to the appropriate packet processing unit PPU 13, as indicated by arrow 122 in Fig. 1. The request includes sending the relevant fields of the packet to the determined packet processing unit PPU 13.

In the following, the packet processing unit PPU 13 processes the received relevant packet fields . As a result, the packet processing unit PPU 13 returns an identification of the determined output port to forward the data packet 112 to, indicated by arrow 124. That is, the packet processing unit PPU 13 returns in the example shown in Fig. 1 the identification LC 14, which means that the data packet 112 needs to be forwarded to line card LC 14.

Other information about the desired packet treatment and alterations, based on the packet processing, may also be sent back to the requesting unit. That information includes, for example, a decision on specific QoS (Quality of Service) treatment of the packet, its mapping onto an MPLS (Multiprotocol Label Switching) label, discarding it or sending it to a control point in case of filtering or splicing with another TCP (Transmission Control Protocol) connection in case of load balancing.

In the next step the line card LC 12 resumes the previously stored packet , applies the desired treatment and alterations to the packet and forwards the data packet 112 to the indicated output, here line card LC 14, denoted by arrow 126. From line card LC 14 the data packet 112 gets fed into the transmission facility connected to the output 114 of line card LC 14 for further transmission as indicated by arrow 128.

TE3020"04388866

Fig. 2 is a high-level block diagram illustrating a preferred embodiment of a router 210 in accordance to the present invention. In Fig. 2, most of the parts shown have equivalents in Fig. 1. Furthermore, the sequence of operation described with reference to Fig. 1 also applies for the embodiment depicted in Fig. 2.

The router 210 comprises a plurality of input/output line cards LC 21 to LC 26, a switching unit 230 and a control unit 232. Each line card LC 21 to LC 26 comprises an input/output port 214 and one packet processing unit PPU 21 to PPU 26, whereby packet processing unit PPU 21 is situated in line card LC 21, packet processing unit PPU 22 is situated in line card LC 22 and so on. Hence, the line cards LC 21 to LC 26 do not only distribute the workload among all other packet processing units PPU 21 to PPU 26, but also among themselves, i.e., in the scheme according to the present invention the packet processing units PPU 21 to PPU 26 are both the requesting units and the processing units.

According to another embodiment in accordance to the present invention the packet processing units PPU 21 to PPU 26 are situated locally at the input ports 214 as part of the line cards LC 21 to LC 26. However, all packet processing units PPU 21 to PPU 26 are still treated as a pool of parallel processing units accessed through the switching unit 230.

In response to a data packet 212 arriving at the input port 214 of line card LC 22 (cf. arrow 220) packet processing unit PPU 23 of line card LC 23 is determined. Again, the data packet 212 consists of a header portion 216 and a data portion 218. As indicated with arrow 222, the packet processing unit PPU 23 of line card LC 23 is then requested to derive a respective output port to forward the data packet 212 to, whereby the respective output port is derived from the piece of packet information, for example, the destination address, associated to the data packet 212. In the following

TE020"0286550

Table 1. Demographic characteristics of the study population	
Age (years)	65.0 ± 10.0
Gender (male/female)	10/10
Education (years)	12.0 ± 2.0
Occupation (white/blue)	10/10
Marital status (married/divorced/widowed)	10/0/0
Smoking status (smoker/nonsmoker)	10/10
Alcohol consumption (yes/no)	10/10
Comorbidities (hypertension/diabetes/cholesterol)	10/10/10
Medication (antihypertensive/antidiabetic/anticholesterol)	10/10/10
Family history (hypertension/diabetes/cholesterol)	10/10/10
Physical activity (yes/no)	10/10
Stress level (high/low)	10/10
Sleep quality (good/poor)	10/10
Depression score (yes/no)	10/10
Overall health status (good/poor)	10/10

[illegible][illegible][illegible][illegible]

the partitioning vector  $p = (p_1, p_2, p_3, \dots, p_M)$ . The partitioning vector  $p$  defines a size of a fraction of the identifier vector space assigned to each packet processing unit PPU 11 to PPU M. The partitioning vector  $p$  is then used to compute a weights vector  $x$  (cf. above, description to Fig. 1) that is used to distribute the load evenly over all packet processing units PPU 31 to PPU L. The partitioning vector  $p$  or the weights vector  $x$  is uploaded to all the line cards LC 31 to LC K as indicated by arrows 336.

As indicated by arrows 334 each packet processing unit PPU 31 to PPU L periodically informs the control unit 332 of its packet processing load, i.e., the number of packets to be processed in an instant of time. The control unit 332 creates from the provided information a load balance vector  $r = (r_1, r_2, r_3, \dots, r_L)$ . In case that the imbalance among the processing units exceeds a certain limit (threshold), the control unit 332 computes a new partitioning vector  $p' = (p'_1, p'_2, p'_3, \dots, p'_L)$  as a function  $p' = g(r, p)$  of the previous partitioning vector  $p$  and the load balance vector  $r$ . Accordingly, a new weights vector  $x$  is calculated as well. Finally, the new partitioning vector  $p'$  or the new weights vector  $x$  is uploaded to all the line cards LC 31 to LC K as indicated by arrows 336.

The present invention can be realized in hardware, software, or a combination of hardware and software. Any kind of computer system - or other apparatus adapted for carrying out the methods described herein - is suited. A typical combination of hardware and software could be a general purpose computer system with a computer program that, when being loaded and executed, controls the computer system such that it carries out the methods described herein. The present invention can also be embedded in a computer program product, which comprises all the features enabling the implementation of the

Computer program means or computer program in the present context mean any expression, in any language, code or notation, of a set of instructions intended to cause a system having an information processing capability to perform a particular function either directly or after either or both of the following a) conversion to another language, code or notation; b) reproduction in a different material form.

[illegible]